

Parallel High Speed Backup for a Storage Area Network (SAN) File System

Abstract of the Disclosure

File system backups are carried out by first generating a list of inodes including associated inode numbers in inode number order that have changed since the last backup operation. A table which has file names and inode numbers for all of the files currently in the file system is also generated. This list is sorted by inode number and the table and list are merged to provide a structure for determining which files are to be backed up. This means that relevant inodes and file names are now provided in a single entity. It is also noted that the structure that results from the merge operation is particularly suitable for being read in blocks which thus permits the backup operation to be carried out in parallel. The task of backing up files is also preferably partitioned by file size or other criteria as opposed to being partitioned simply by the number of files assigned to be backed up by any one processor in a distributed or parallel data processing system.

15